Curve-Sensitive Cuttings*

Vladlen Koltun† Micha Sharir‡

December 16, 2004

Abstract

We introduce $(1/r)$-cuttings for collections of surfaces in 3-space, such that the cuttings are sensitive to an additional collection of curves. Specifically, let S be a set of n surfaces and let C be a set of m curves in $\mathbb{R}^3$, all of constant description complexity. Let $1 \le r \le \min\{m, n\}$ be a given parameter. We show the existence of a $(1/r)$-cutting Ξ of S of size $O(r^{3+\varepsilon})$, for any $\varepsilon > 0$, such that the number of crossings between the curves of C and the cells of Ξ is $O(mr^{1+\varepsilon})$. The latter bound improves, by roughly a factor of r, the bound that can be obtained for cuttings based on vertical decompositions. We view curve-sensitive cuttings as a powerful tool for various scenarios that involve curves and surfaces in three dimensions. As a preliminary application, we use the construction to obtain a bound of $O(m^{1/2}n^{2+\varepsilon})$, for any $\varepsilon > 0$, on the complexity of the multiple zone of m curves in the arrangement of n surfaces in 3-space. After the conference publication of this paper [15], curve-sensitive cuttings were applied to derive an algorithm for efficiently counting triple intersections among planar convex objects in three dimensions [12], and we expect additional applications to arise in the future.

1 Introduction

Motivation. $(1/r)$-cuttings (see below for definitions) have attracted considerable attention in the computational geometry community, as they turned out to be crucial to the solution of many central problems in the field [5, 6, 7, 8, 9, 10, 14, 16, 17]. For some applications, special properties possessed by the cutting can lead to improved results. For instance, the tree structure of *hierarchical cuttings* [6] is of great help in numerous settings [4, 17].

We construct a $(1/r)$-cutting for a collection of surfaces in 3-space, such that the cutting is sensitive, in the sense defined below, to a collection of curves given as additional input to the construction. We apply this cutting to obtain a bound of $O(m^{1/2}n^{2+\varepsilon})$, for any $\varepsilon > 0$,

*Work on this paper was initiated at the DIMACS Workshop on Geometric Graph Theory, held at Rutgers University, Fall 2002. This work has been supported by a grant from the U.S.-Israeli Binational Science Foundation, by a grant from the Israel Science Fund (for a Center of Excellence in Geometric Computing), by NSF Grants CCR-97-32101, CCR-00-98246, and CCR-01-21555 and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University.

†Computer Science Division, University of California, Berkeley, CA 94720-1776, USA; *vladlen@cs.berkeley.edu*.

‡School of Computer Science, Tel Aviv University, Tel-Aviv 69978, Israel, and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA; *michas@post.tau.ac.il*.

on the complexity of the multiple zone of m curves in the arrangement of n surfaces in 3-space, all of constant description complexity. The multiple zone is defined as the collection of all cells of the arrangement of the given surfaces that are crossed by at least one of the curves. It is a generalization of both the concept of the zone of a curve in an arrangement [3, 13] and the widely studied notion of many faces/cells in arrangements [2].

We expect curve-sensitive cuttings to find additional uses in contexts that involve the interaction of curves and surfaces. It has already been applied, after the conference publication of this paper [15], to derive an algorithm for efficiently counting triple intersections among planar convex objects in three dimensions [12].

Overview. Let S be a set of n surfaces in $\mathbb{R}^3$ of constant description complexity, and let C be a set of m curves in $\mathbb{R}^3$ of constant description complexity; that is, each surface and curve is defined as a Boolean combination of a constant number of polynomial equations and inequalities of constant maximum degree. Let $1 \le r \le \min\{m, n\}$ be a given parameter. A $(1/r)$-*cutting* of S is a subdivision of 3-space into connected cells, each of constant description complexity, so that each cell is crossed by at most n/r surfaces of S. We wish to construct a $(1/r)$-cutting Ξ of S of size near $O(r^3)$, so that the number of pairs (c, τ), where $c \in C$, τ a cell of Ξ, and $c \cap \tau \neq \emptyset$, is near $O(mr)$; that is, the average number of cells of Ξ crossed by a curve of C is near $O(r)$.

A standard method (in fact, the only general-purpose method known to date) for constructing a $(1/r)$-cutting for arrangements of non-linear surfaces is to take an appropriate random sample R of the surfaces of S, and to construct the *vertical decomposition* of the arrangement $\mathcal{A}(R)$ of R [18]. The construction of this decomposition proceeds in two stages. First, for every edge of $\mathcal{A}(R)$ and every vertical tangency curve (also known as the *silhouette*) on every surface of R, we erect a 2-dimensional vertical *visibility wall*, defined as the union of all z-vertical segments that have an endpoint on this edge (or curve) and are interior-disjoint from all surfaces of R. This first stage results in a decomposition of $\mathcal{A}(R)$ into vertical pseudo-prisms, such that the floor of each prism, if it exists, is contained in a single surface of R, and similarly for the ceiling of each prism. However, the combinatorial complexity of a single prism can still be fairly high.

In the second stage of the construction we refine the decomposition as follows. For every prism as above, consider its projection onto the xy-plane. It is a 2-dimensional semi-algebraic set, which we decompose in the plane by erecting zero, one, or two y-vertical (possibly infinite) visibility segments on each of its vertices and y-vertical tangency points on its edges, where a visibility segment is defined as a maximal y-vertical segment that has an end-point on this vertex (or tangency point), is contained in the considered prism projection, and is interior-disjoint from its boundary. We then erect z-vertical 2-dimensional walls inside the original prism, defined as its intersection with the z-vertical walls spanned by all the y-vertical segments erected by the planar decomposition. Repeating this process for each of the above prisms decomposes $\mathcal{A}(R)$ into cells of constant description complexity.

We can choose R as a single sample from S of size $ar \log r$, for an appropriate absolute constant a. It can then be argued that, with high probability, the resulting vertical decomposition of $\mathcal{A}(R)$ is indeed a $(1/r)$-cutting. This is a consequence of the probabilistic analyses of Haussler and Welzl [14] and of Clarkson [9]. Using a variant of the method of Chazelle and Friedman [7] or of Chazelle [6], slightly reduces the size of the resulting cutting from $O(r^3 \log^3 r)$ to $O(r^3)$.

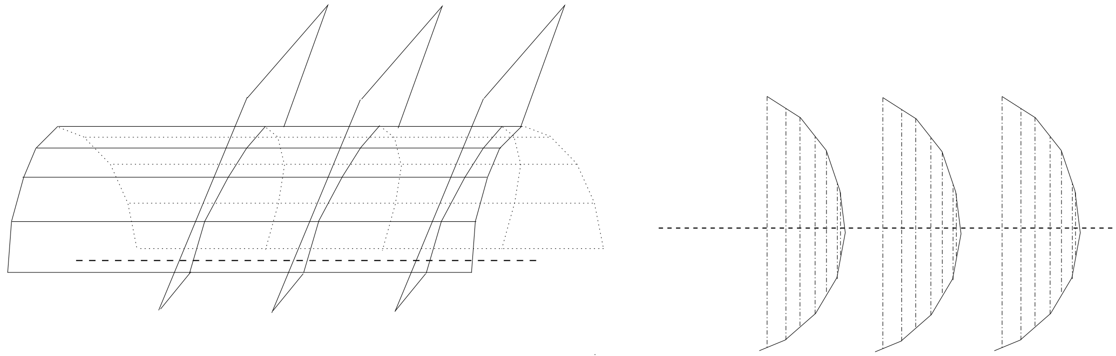

Figure 1: A curve (the x-axis, shown dashed) crossing a quadratic number of cells of the vertical decomposition. (a) A side view of the input set. (b) A view from above of the second-step subdivision of the cells mentioned in the text.

Unfortunately, vertical decompositions may fail to satisfy our requirement concerning the number of crossings between the curves of C and the cells of the cutting. In fact, a curve may cross nearly $\Omega(r^2)$ such cells. An example is shown in Figure 1, where R is a collection of r planes. Half of them are parallel to the x-axis and pass above it, all appearing on the lower envelope of this group, which looks like a tunnel in the x-direction with a convex roof that is symmetric about the xz-plane. The remaining $r/2$ planes are all parallel to the y-axis, and form a fixed angle, say $45°$, with the xy-plane. These latter planes are sufficiently separated from each other, so that their portions that lie above the xy-plane and below the lower envelope of the first group have pairwise disjoint xy-projections. The x-axis crosses $\Theta(r^2)$ cells of the vertical decomposition of these planes: Indeed, the first decomposition step creates (among others) $r/2$ cells whose top facet is the portion of some slanted plane of the second group that lies below the lower envelope of the first group. The second decomposition step subdivides each of these cells into $\Theta(r)$ subcells, and the x-axis crosses them all.

In contrast, the *undecomposed* arrangement of $\Theta(r \log r)$ surfaces is sensitive to the curves of C, because each curve crosses each surface at $O(1)$ points, so it crosses $O(r \log r)$ cells of the arrangement. However, the undecomposed arrangement is generally not a $(1/r)$-cutting. On the other hand, the decomposed arrangement is (with high probability) a $(1/r)$-cutting, but, as we have just seen, it may fail to be sensitive to C. (Actually, as we will show in Section 2.1, the first stage of the vertical decomposition is also sensitive to C, but in general it is still not a $(1/r)$-cutting.)

In this paper we describe a technique that achieves the better of both worlds and constructs cuttings that satisfy the desired properties. The construction proceeds by taking a sample R of the surfaces, as described above, and decomposing $\mathcal{A}(R)$ into vertical prisms using the *first stage* of the vertical decomposition construction. Inside each prism we construct a decomposition that takes into account the parts of the curves of C that lie inside the prism. Specifically, we construct a hierarchical sequence of cuttings, somewhat reminiscent of the construction in Chazelle [6], that reduces the number of crossings between the curves of C and the boundaries of the cells of the cuttings. We are able to guarantee that the curves of C are not cut more than $O(mr^{1+\varepsilon})$ times, for any $\varepsilon > 0$, overall.

Before describing our results in detail, we remark that we can construct an alternative (and simpler) curve-sensitive decomposition scheme for the special case where the surfaces are planes and the curves are lines (as in the example of Figure 1), by using the Dobkin-Kirkpatrick hierarchical decomposition [11] in each cell of $\mathcal{A}(R)$. This approach, however, does not extend to general curves and surfaces. (An expanded discussion of this remark is given in the application paper [12].)

2 A Curve-Sensitive Decomposition

In this section we present a new decomposition scheme that is a $(1/r)$-cutting for S and satisfies the desired bounds on the number of cells and on the number of curve-cell crossings. For simplicity of exposition, we will base our analysis on a single random sample of surfaces from S (rather than the more elaborate repeated-sampling scheme of [7]). Moreover, we consider samples of size r (rather than $\Theta(r \log r)$). This simplifies the calculations, but will only produce a $O(\log r/r)$-cutting. We get the desired cutting by simply replacing r, at the end of the analysis, by the above larger sample size.

2.1 First Stage of the Decomposition

We begin with taking a random sample R of r surfaces of S, and a random sample R' of r curves of C. We form the arrangement $\mathcal{A}(R)$ of R, and apply to it the first step of the vertical decomposition. That is, we erect vertical walls up and down from each curve of intersection of pairs of surfaces in R, as well as from the silhouette of each surface in R; the walls are extended until they hit another surface of R, or, failing that, all the way to $\pm\infty$. In addition, we erect similar vertical walls from each curve $c \in R'$, which are also extended to the first surface above and below.

Let $\mathcal{A}_1 = \mathcal{A}_1(R, R')$ denote the resulting decomposition. Note that each cell τ of $\mathcal{A}_1$ is a vertical prism-like cell: the intersection of each vertical line with τ is connected. However, the xy-projection τ^* of τ can have arbitrary shape and complexity.

For each cell τ of $\mathcal{A}_1$, let ξ_τ denote its combinatorial complexity (i.e., the number of vertices, edges and faces on its boundary), and let C_τ denote the set of all connected components of the nonempty intersections between τ and the curves of C. Let $\lambda_q(r)$ denote, as usual, the maximum length of a Davenport-Schinzel sequence of order q on r symbols [18], and put $\beta_q(r) = \lambda_q(r)/r$, which is thus an extremely slow-growing function of r. We have

Lemma 2.1. *(a) The number of cells of $\mathcal{A}_1$ and their overall combinatorial complexity are both $O(r^3\beta_q(r))$, for an appropriate parameter q that depends on the algebraic complexity of the curves of C and the surfaces of S.*

(b) $\sum_{\tau \in \mathcal{A}_1} |C_\tau| = O(mr\beta_q(r))$.

Proof: Let γ be a fixed curve, which is either a curve in C, or an intersection curve of two surfaces in R, or the silhouette of a surface in R. Let V_γ denote the vertical 2-manifold (wall) spanned by γ. Let V_γ^+ (resp., V_γ^-) denote the portion of V_γ that lies above (resp., below) γ. Let $\mathcal{A}^+$ (resp., $\mathcal{A}^-$) denote the cross section of $\mathcal{A}(R)$ with V_γ^+ (resp., V_γ^-). By construction, any point at which γ crosses the boundary of some cell of $\mathcal{A}_1$ must either be the vertical projection on γ of a vertex of the lower envelope of $\mathcal{A}^+$, or a vertex of the upper envelope of $\mathcal{A}^-$ (or of both, if the vertex lies on γ itself), or a point that lies vertically above or below a point on another curve of R' (so that the two points are *vertically visible* in $\mathcal{A}(R)$). The complexity of each envelope is $O(\lambda_q(r)) = O(r\beta_q(r))$, for an appropriate constant q [18], and the number of times γ passes above or below any curve of R' is $O(r)$ (over all curves of R'). This readily implies the lemma: Part (b) is an immediate consequence, while part (a) follows by applying this bound to each of the $O(r^2)$ intersection and $O(r)$ silhouette curves arising in the sample. □

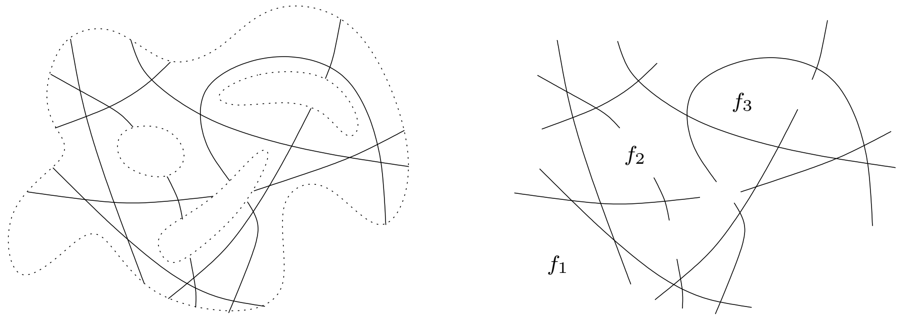

(a) (b)

Figure 2: Stage 2 of the decomposition. (a) The curves of $Q \subseteq C^*_{\tau_0}$ (solid) and $\partial\tau^*$ (dotted). (b) The external faces of $\mathcal{A}(Q)$; note that f_2 contains two components of $\partial\tau^*$.

2.2 Second Stage of the Decomposition

After constructing the decomposition $\mathcal{A}_1$, we perform a second decomposition step, which decomposes each cell τ of $\mathcal{A}_1$ as follows. Let $\partial\tau^*$ denote the boundary of τ^* and let h_τ denote the number of internal boundary components ("holes") of τ^*. Note that $h_\tau \leq \xi_\tau$. Since τ^* need not be simply connected, $\partial\tau^*$ may consist of more than one connected component. (I.e., h_τ may be strictly positive.) The potential existence of many components of $\partial\tau^*$ is the main source of technical difficulty in the analysis of our decomposition.

Put $m_\tau = |C_\tau|$. Let C^*_τ denote the set of the xy-projections of the arcs of C_τ. Let X_τ denote the number of intersections between the curves of C^*_τ. This is also equal to the number of *vertical visibility segments* between pairs of curves of C_τ, where such a segment is parallel to the z-axis and connects a point on one curve to a point on the other (and is thus fully contained in τ). We clearly have

$$\sum_{\tau \in \mathcal{A}_1} X_\tau = O(m^2). \qquad (1)$$

In what follows, through the bulk of this section, we assume that $X_\tau \geq m_\tau$. The alternative case $X_\tau < m_\tau$ is considerably simpler to handle and will be described later in the analysis.

If $X_\tau = o(m_\tau^2)$, we carry out a preliminary decomposition stage that covers τ^* by the union of simpler-shaped subcells, so that, within each such subcell τ_0, the number of intersections between the curves of C^*_τ that cross τ_0 is roughly the square of the number of such curves. We employ a standard approach that proceeds as follows. (See, e.g., [5].) Put $s = s_\tau = \lceil m_\tau^2/X_\tau \rceil$. We distinguish between the following two cases.

(a) Suppose first that $s \le \xi_\tau$. We sample each curve of C^*_τ with probability s/m_τ. This produces a random sample R'' of expected size s. The expected complexity of $\mathcal{A}(R'')$ is $O(s + (s/m_\tau)^2 X_\tau) = O(s)$, since each intersection counted in X_τ becomes a vertex of $\mathcal{A}(R'')$ with probability $(s/m_\tau)^2$. We construct the vertical decomposition of $\mathcal{A}(R'')$, and argue that, with high probability, it consists of $O(s)$ trapezoids, each of which is crossed by at most $O((m_\tau/s)\log s)$ curves of C^*_τ. We next apply a modified version of the analysis of Chazelle and Friedman [7] to refine the decomposition, so that each of its cells is crossed by at most m_τ/s curves of C^*_τ, while the number of cells remain $O(s)$.

Since the setup here is somewhat different from that in [7], we present details of the construction and of its analysis. This is done as follows. We take each cell Δ of the vertical decomposition that is crossed by tm_τ/s curves of C^*_τ, for any $t > 1$, draw a random sample R''_Δ of $ct\log t$ of these curves, for an appropriate sufficiently large constant c, construct the vertical decomposition of the arrangement $\mathcal{A}(R''_\Delta)$, and clip each resulting cell to Δ. With high probability, each cell in the resulting decomposition is crossed by at most m_τ/s curves of C^*_τ, provided c is chosen sufficiently large. To estimate the overall number of cells, we apply Lemma 2.2 of Agarwal et al. [1] which, in our context, asserts that the expected number of cells that are crossed by at least jm_τ/s curves is $O(2^{-j})$ times the expected number of cells in a random sample of s/j curves of C^*_τ. The latter expected number of cells is easily seen to be $O(s/j)$, and thus the overall expected number of new cells is

$$O(s) + \sum_{j\ge 1} O(2^{-j}s/j) = O(s),$$

as claimed.

As we do throughout the analysis, we assume that all those subsamples meet their expected values, so that this property holds with certainty. This assumption can be made effective, e.g., by resampling at each stage of the construction until a good sample is obtained. See a remark to that effect following Theorem 2.3.

These trapezoids are the cells τ_0 of the cutting decomposition (or, rather, covering) of τ^*. Each cell τ_0 contains on average $X_\tau/s = O(X_\tau^2/m_\tau^2)$ crossings between curves of C^*_τ, which is roughly the square of the number $O(m_\tau/s) = O(X_\tau/m_\tau)$ of these curves that cross τ_0. It is important to note that this decomposition is defined only in terms of the curves in C^*_τ, and is thus not necessarily confined to within τ^*. Thus our trapezoids constitute a *covering* of τ^*. (Nevertheless, since all the curves of C^*_τ are fully contained in τ^*, the portion of the covering outside τ^* is uninteresting; it is constructed simply because we do not want at this stage to let $\partial\tau^*$ affect the construction.) We shall later, towards the end of this section, take care to clip the new cells to within τ^*.

(b) Suppose next that $s > \xi_\tau$. We then sample each curve of C^*_τ with probability ξ_τ/m_τ. Note that this quantity is indeed at most 1, because $s \le m_\tau$ (which follows from the assumption $X_\tau \ge m_\tau$) and $\xi_\tau < s$. This produces a random sample R'' of expected size ξ_τ. The expected complexity of $\mathcal{A}(R'')$ is $O(\xi_\tau + (\xi_\tau/m_\tau)^2 X_\tau) = O(\xi_\tau)$, since $\xi_\tau < s$. We apply the same decomposition construction as in the preceding case, obtaining a new collection of $O(\xi_\tau)$ trapezoids, each of which is crossed by at most m_τ/ξ_τ curves of C^*_τ. These trapezoids are the cells τ_0 of the cutting-cover of τ^*.

This concludes the description of the preliminary covering of τ^* that is constructed only if $X_\tau = o(m_\tau^2)$. If $X_\tau = \Theta(m_\tau^2)$, we have $s = O(1)$ and the first case applies; we cover τ by a single τ_0, which we take to be the entire xy-plane.

We now apply an additional decomposition step to each cell τ_0 of this preliminary cutting. This decomposition consists of a recursively constructed hierarchical sequence of cuttings of the subset $C_{\tau_0}^*$ of those curves of C_τ^* that cross τ_0, clipped to within τ_0. This decomposition is somewhat reminiscent of the hierarchical cutting construction of Chazelle [6]. We begin by choosing a sufficiently large constant ρ, to be used throughout the construction. Put $m_{\tau_0} = |C_{\tau_0}^*|$.

First level in the hierarchy. We draw a random sample Q of ρ arcs of $C_{\tau_0}^*$ and consider all the faces of the planar arrangement $\mathcal{A}(Q)$ that contain components of $\partial\tau^*$. By the definition of C_τ, the arcs of $C_{\tau_0}^*$ are contained within τ^*, and thus each component of $\partial\tau^*$ lives in a single (not necessarily distinct) face of $\mathcal{A}(Q)$. We refer to such faces as the *external* faces of $\mathcal{A}(Q)$. Note also that, as defined, those faces are not confined to within τ_0 nor within τ^*. That is, $\partial\tau^*$ is not part of $\mathcal{A}(Q)$ and does not delimit any face of it. However, each component γ of $\partial\tau^*$ bounds a connected component of the complement of τ^* which is fully disjoint from all the arcs of Q (or of $C_{\tau_0}^*$ for that matter). See Figure 2.

For each external face f of $\mathcal{A}(Q)$, we compute the 2-dimensional vertical decomposition of f into vertical pseudo-trapezoids (see, e.g., [18]), which we refer to as *trapezoids* or *subcells*. With high probability (greater than, say, $1 - 1/\rho$), each resulting subcell σ is crossed by at most $\frac{am_{\tau_0}}{\rho}\log\rho$ curves of $C_{\tau_0}^*$, for an appropriate absolute constant a [9, 14]. As above, we assume that Q is a sample that satisfies this property. For each connected component γ of $\partial\tau^*$, the face f_γ of $\mathcal{A}(Q)$ that contains γ consists of $O(\rho\beta_q(\rho))$ subcells [18], so the total number of crossings between the arcs of $C_{\tau_0}^*$ and these subcells is $O(m_{\tau_0}\beta_q(\rho)\log\rho)$. Let κ_{τ_0} denote the number of distinct external faces of $\mathcal{A}(Q)$. Then we get a total of $O(\kappa_{\tau_0}\rho\beta_q(\rho))$ external trapezoids,[1] and the total number of crossings between the arcs of $C_{\tau_0}^*$ and these subcells is $O(\kappa_{\tau_0}m_{\tau_0}\beta_q(\rho)\log\rho)$.

An obvious upper bound on κ_{τ_0} is $1 + h_{\tau_0}$, where h_{τ_0} denotes the number of internal connected components of $\partial\tau^*$ that are fully contained in τ_0 (boundary components that cross $\partial\tau_0$ all lie in the single unbounded face of $\mathcal{A}(Q)$), but we will use in the following analysis a more refined bound. The need for a refined analysis comes from the observation that, at this initial stage of the hierarchy, the total number of faces of $\mathcal{A}(Q)$ is only a constant (at most $O(\rho^2)$), whereas h_{τ_0} can be much larger. Note that, trivially,

$$\sum_{\tau_0} h_{\tau_0} \le h_\tau \le \xi_\tau. \tag{2}$$

We also have $h_\tau = O(r)$, because we can charge each internal component of $\partial\tau^*$ either to a complete connected component of an intersection curve between the surface of R forming the floor of τ with another surface in R, or to a similar intersection component involving the surface forming the ceiling of τ, or to a complete connected component of the silhouette of some surface of R (which is completely contained in the interior of τ), and the overall number of such components is clearly $O(r)$. In fact, applying this analysis to all the cells τ

[1] The number of external trapezoids is proportional to the combined complexity of the external faces. In general, better bounds are known for the complexity of κ_{τ_0} faces in an arrangement of ρ curves (see, e.g., [8]), but the cruder bound that we use suffices for our purposes.

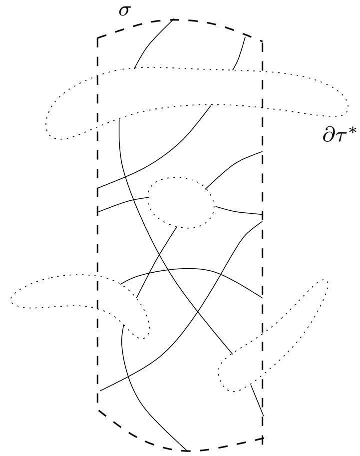

Figure 3: An external trapezoid σ (dashed), the portions of $\partial\tau^*$ that meet σ (dotted), and the arcs in C^*_σ (solid).

of $\mathcal{A}_1$ together, we obtain the following bound, which is crucial for our analysis.

$$\sum_{\tau \in \mathcal{A}_1} h_\tau = O(r^2). \tag{3}$$

In addition to decomposing the external faces as described above, we also partition the remainder of $\mathcal{A}(Q)$ (its *internal portion*) into vertical trapezoids. In doing so, we erase all the edges of $\mathcal{A}(Q)$ that are contained in the interior of the internal portion, and retain only the edges that also bound the external faces. Thus the number of trapezoids into which the internal portion is partitioned is also $O(\kappa_{\tau_0}\rho\beta_q(\rho))$. The total number of crossings between the arcs of $C^*_{\tau_0}$ and these internal subcells is $O(\kappa_{\tau_0} m_{\tau_0}\rho\beta_q(\rho))$. (Here we can no longer claim that each internal trapezoid is crossed by only a small number of curves, because it is not necessarily disjoint from the sampled curves in Q, so this bound is larger than the bound claimed for external trapezoids, by nearly a factor of ρ.)

Second level in the hierarchy. We now apply a second partitioning step[2] within each external trapezoid σ that has a nonempty intersection with $\partial\tau^*$. (All other external and internal trapezoids are not decomposed any further.) Let C^*_σ denote the set of connected components of the intersections of the curves in $C^*_{\tau_0}$ with σ. As in the preceding step, σ is not necessarily contained in τ^*; however, each arc in C^*_σ lies fully in $\sigma \cap \tau^*$. See Figure 3.

We draw a random sample Q_σ of ρ curves of C^*_σ, and compute all the faces of the planar arrangement $\mathcal{A}(Q_\sigma)$ that contain components of $\partial\tau^*$. As above, each component of $\partial\tau^*$ lives in a single ('external') face of $\mathcal{A}(Q_\sigma)$. Again, those faces are not necessarily distinct. This time, however, all external faces, with the exception of the unbounded one, are confined to within σ. Boundary components γ of $\partial\tau^*$ that intersect σ are of two types: those that are fully contained in the interior of σ, and those that cross $\partial\sigma$. All components γ of the second type lie in the same (unbounded) face of $\mathcal{A}(Q_\sigma)$. Let h_σ, κ_σ denote, respectively,

[2]To help the reader follow the construction, we present the second stage explicitly and separately, even though it is a special case of the general recursive step, described later. As a matter of fact, it is also similar to the first-level partitioning.

the number of components γ of the first type, and the number of distinct external faces of $\mathcal{A}(Q_\sigma)$. Clearly, $\kappa_\sigma \le 1 + h_\sigma$, and $\sum_\sigma h_\sigma \le h_\tau$ (where the sum extends over all σ and all τ_0). Again, however, we will have to use a more refined bound for κ_σ in what follows.

For each external face f of $\mathcal{A}(Q_\sigma)$, we compute the 2-dimensional vertical decomposition of f. With high probability (larger than $1 - 1/\rho$), each resulting subcell σ' is crossed by at most

$$\left(\frac{a \log \rho}{\rho}\right)^2 m_{\tau_0}$$

curves of C^*_σ, and, as above, we assume that Q_σ is a sample that does satisfy this property. For each connected component γ of $\partial\tau^*$ that meets σ, the face f_γ of $\mathcal{A}(Q_\sigma)$ that contains γ consists of $O(\rho\beta_q(\rho))$ subcells. Summing over all boundary components of $\partial\tau^*$ that meet σ, we get a total of $O(\kappa_\sigma\rho\beta_q(\rho))$ external trapezoids, and the total number of crossings between the arcs of C^*_σ and these subcells is

$$O(\kappa_\sigma m_{\tau_0} \beta_q(\rho) \log^2 \rho/\rho).$$

Summing these bounds over all external trapezoids σ, we obtain bounds for the overall number of external trapezoids in the second hierarchical partitioning step, and the total number of crossings between arcs in $C^*_{\tau_0}$ and these trapezoids. These bounds are, respectively,

$$\sum_\sigma O(\kappa_\sigma \rho \beta_q(\rho)), \tag{4}$$

and

$$\sum_\sigma O(\kappa_\sigma m_{\tau_0} \beta_q(\rho) \log^2 \rho/\rho),$$

where these sums are over all external trapezoids σ in $\mathcal{A}(Q)$.

As above, we also partition the remainder internal portions of the arrangements $\mathcal{A}(Q_\sigma)$, over all trapezoids σ, into vertical trapezoids, using, as above, only the edges and vertices of these internal portions that bound also the external portions. Thus, the overall number of internal trapezoids is also bounded by (4), and the total number of crossings between arcs in $C^*_{\tau_0}$ and these internal trapezoids is at most

$$\sum_{\sigma \text{ an external trapezoid in } \mathcal{A}(Q)} O(\kappa_\sigma m_{\tau_0} \beta_q(\rho) \log \rho).$$

Recursive construction of the hierarchy. The above process is repeated recursively, each recursion stage refining the decomposition inside those "external" trapezoids constructed in the previous stage that are still crossed by (or contain) boundary components of $\partial\tau^*$. Let $j = j_{\tau_0}$ be the smallest integer such that

$$\rho^j \ge \xi_\tau / s.$$

We stop the recursive decomposition process after j steps. In particular, if $\xi_\tau < s$, there is no recursion, and τ_0 remains intact. Otherwise, we have $\rho^j = \Theta(\xi_\tau/s)$. Let us for now only consider the (much more involved) case $\xi_\tau \ge s$.

By an appropriate extension of the preceding arguments, the overall number of external and internal trapezoids produced in the i-th step, for any $i = 1, \ldots, j$, is at most

$$\sum_{\sigma \text{ an external trapezoid in some } \mathcal{A}(Q_{\sigma'})} O(\kappa_\sigma \rho \beta_q(\rho)), \tag{5}$$

where σ' is an external trapezoid constructed in the preceding $(i-1)$-st step which intersects $\partial\tau^*$. With high probability (which we turn into certainty by choosing "good" samples $Q_{\sigma'}$), each external trapezoid constructed at the i-th step is crossed by at most

$$O\left(\left(\frac{a \log \rho}{\rho}\right)^i m_{\tau_0}\right)$$

curves of C_τ^*, and each such internal trapezoid is crossed by at most

$$O\left(\left(\frac{a \log \rho}{\rho}\right)^{i-1} m_{\tau_0}\right)$$

curves. Hence, the number of crossings between the arcs of $C_{\tau_0}^*$ and the external trapezoids is at most

$$\sum_\sigma O\left(\kappa_\sigma m_{\tau_0} \beta_q(\rho) \frac{a^i \log^i \rho}{\rho^{i-1}}\right),$$

and the number of crossings between the arcs of $C_{\tau_0}^*$ and the internal trapezoids is at most

$$\sum_\sigma O\left(\kappa_\sigma m_{\tau_0} \beta_q(\rho) \frac{a^{i-1} \log^{i-1} \rho}{\rho^{i-2}}\right), \tag{6}$$

where these sums are over all external trapezoids σ in some $\mathcal{A}(Q_{\sigma'})$.

Bounding the number of trapezoids. We continue to assume in what follows that $\xi_\tau \geq s$; otherwise τ_0 remains a single trapezoid. Let us analyze the number of trapezoids in more detail. Let γ be a boundary component of $\partial\tau^*$. If at some step i, γ crosses the boundary of some external trapezoid(s), it has no effect on the quantities κ_σ from this step on (inclusive). If on the other hand γ remains confined to the interior of a single external trapezoid σ, then it may add 1 to κ_σ, but it will not affect $\kappa_{\sigma'}$, for any other external trapezoid σ' produced at this step.

Elaborating this observation, we consider the tree $\mathcal{T}$ of all external trapezoids as they are generated during the recursive process. The root of the tree is τ_0, and the children of each external trapezoid σ are the external trapezoids that are constructed in the decomposition of σ. We say that a trapezoid is *pregnant* if it completely contains a component of $\partial\tau^*$ in its interior. Otherwise it is *empty*. An empty trapezoid can spawn at most $c\rho\beta_q(\rho)$ sub-trapezoids (in a single decomposition step), for some constant c, whereas a pregnant trapezoid containing t components of $\partial\tau^*$ in its interior can spawn as many as $(t+1)c\rho\beta_q(\rho)$ sub-trapezoids, but no more than $c\rho^2$, which is the maximum number of trapezoids that can be generated in a single decomposition step (for simplicity, we use the same constant c in both bounds). Moreover, the total number of pregnant trapezoids, over the entire tree, is only $O(h_{\tau_0})$.

The empty trapezoids are organized in subtrees, each rooted at some pregnant trapezoid. (If there are no pregnant trapezoids, the empty trapezoids comprise the entire tree $\mathcal{T}$, and the analysis becomes considerably simpler.) Consider such a subtree rooted at a pregnant trapezoid at depth i (where the root of $\mathcal{T}$ is at depth 0). The degree of each node in the subtree is at most $c\rho\beta_q(\rho)$, so the size (or, more precisely, the number of leaves) of the subtree is at most $(c\rho\beta_q(\rho))^{j-i}$ (recall that j is the depth of the entire recursion). We choose a threshold depth k, and distinguish between the cases $i \le k$ and $i > k$. In the former case, the total number of trees whose roots are at depth i is at most $c^i\rho^{2i}$, and their total size (i.e., number of leaves) is thus at most

$$c^i\rho^{2i} \cdot (c\rho\beta_q(\rho))^{j-i} = c^j\rho^{i+j}(\beta_q(\rho))^{j-i}.$$

Summing these bounds over all depths $i = 0, \ldots, k$, we obtain a total size of

$$O(c^j\rho^{k+j}(\beta_q(\rho))^{j-k})$$

trapezoids.

In the latter case ($i > k$), we bound the total number of trees whose roots are at depth greater than k simply by $O(h_{\tau_0})$, and bound the size of any such subtree by $(c\rho\beta_q(\rho))^{j-k}$. Hence, the total size of these subtrees is at most

$$O(h_{\tau_0}) \cdot (c\rho\beta_q(\rho))^{j-k}.$$

To fix the value of k, we first assume that $(c\rho^2)^j \le h_{\tau_0}$, set $k = j$, and note that only the case $i \le k$ remains relevant. The overall number of external trapezoids produced within τ_0 under this assumption is

$$O(c^j\rho^{2j}) = O(h_{\tau_0}).$$

Assuming now that $(c\rho^2)^j > h_{\tau_0}$, we choose k so that $(c\rho^2)^k = \Theta(h_{\tau_0})$, and assume that ρ is a sufficiently large constant, as a function of a prescribed $\varepsilon > 0$, to conclude that the overall number of external trapezoids in this setting is

$$O(h_{\tau_0}^{1/2}\rho^{j(1+\varepsilon)}) = O\left(h_{\tau_0}^{1/2}\left(\frac{\xi_\tau}{s}\right)^{1+\varepsilon}\right).$$

By construction, the number of internal trapezoids has the same asymptotic upper bound.

Note that when $h_{\tau_0} = 0$, there is only one subtree, with $(c\rho\beta_q(\rho))^j = O((\xi_\tau/s)^{1+\varepsilon})$ trapezoids.

We sum the above two bounds over all cells τ_0 (for the fixed first-stage cell τ), use the Cauchy-Schwarz inequality, and the facts that $\sum_{\tau_0} h_{\tau_0} \le h_\tau$ and that the number of trapezoids τ_0 is $O(s)$, and cater to both cases $h_{\tau_0} > 0$ and $h_{\tau_0} = 0$, to conclude that the total number of trapezoids into which τ is partitioned is

$$O\left(\sum_{\tau_0} h_{\tau_0}\right) + O\left(\left(\frac{\xi_\tau}{s}\right)^{1+\varepsilon}\right) \cdot \sum_{\tau_0} \max\{1, h_{\tau_0}\}^{1/2} =$$
$$O\left(\sum_{\tau_0} h_{\tau_0}\right) + O\left(\left(\frac{\xi_\tau}{s}\right)^{1+\varepsilon}(h_\tau + s)^{1/2}s^{1/2}\right) =$$
$$O\left(h_\tau + (1 + h_\tau^{1/2})\xi_\tau^{1+\varepsilon}\right). \quad (7)$$

We now cater to the case $\xi_\tau < s$. In this case, τ is covered by $O(\xi_\tau)$ trapezoids τ_0, and each of them remains intact, so the total number of trapezoids is $O(\xi_\tau)$, which is subsumed in the bound (7).

Bounding the number of curve-cell crossings. Next consider the bounds (6) on the number of curve-cell crossings, and analyze them in more detail, using our tree representation of the external trapezoids. We continue to assume that $X_\tau \geq m_\tau$, and that $\xi_\tau \geq s$. For simplicity of exposition, we only consider crossings with external trapezoids, observing that at each step of the construction, the number of internal trapezoids has the same upper bound as the number of external trapezoids, and that, with high probability (which, as usual, we take to hold with certainty), the bound on the number of curves of C^*_σ that cross an internal trapezoid is at most $\rho/(a \log \rho)$ times larger than the same bound for external trapezoids. Hence, up to this constant, the number of crossings with internal trapezoids has the same upper bound as the number of crossings with external trapezoids, so we only concentrate on bounding the latter quantity.

Consider our tree $\mathcal{T}$ of external trapezoids. As in the preceding analysis, we distinguish between the cases $h_{\tau_0} > 0$ and $h_{\tau_0} = 0$. We only treat the case $h_{\tau_0} > 0$; the other case is handled similarly, by replacing h_{τ_0} by 1. With high probability (which, as usual, we take to hold with certainty), an external trapezoid at depth i is crossed by at most $\left(\frac{a \log \rho}{\rho}\right)^i m_{\tau_0}$ curves of $C^*_{\tau_0}$. We fix a threshold value k as above, taking also into consideration the case where $(c\rho^2)^j \leq h_{\tau_0}$. Suppose first that $i \leq k$. The number of external trapezoids at depth i is at most $c^i \rho^{2i}$, so the overall number of curve-cell crossings with these trapezoids is at most $(ac\rho \log \rho)^i m_{\tau_0}$. Summing this over all depths $i = 0, \ldots, k$, we get a total of $O((ac\rho \log \rho)^k m_{\tau_0})$ crossings. For both possible values of k, by the choices of j and ρ, the above bound can be written as

$$O\left((ac^{1/2} \log \rho)^k (c^{1/2}\rho)^k m_{\tau_0}\right) = O\left(h_{\tau_0}^{1/2} (\xi_\tau / s)^\varepsilon m_{\tau_0}\right).$$

Consider next the case $i > k$ (which only applies when $(c\rho^2)^j > h_{\tau_0}$). The number of external trapezoids at depth i can be estimated as follows. All these trapezoids belong to $O(h_{\tau_0})$ subtrees rooted at the pregnant trapezoids, or, if $h_{\tau_0} = 0$, to the entire tree $\mathcal{T}$. To maximize the number of our trapezoids, the subtrees should be rooted as close to the root of $\mathcal{T}$ as possible. By the choice of k, it is easily seen that this happens when all the pregnant nodes lie roughly at level k of $\mathcal{T}$. Assuming this "worst-case" scenario, the number of external trapezoids at depth i is at most

$$O\left((c\rho^2)^k \cdot (c\rho\beta_q(\rho))^{i-k}\right) = O\left(c^i \rho^{i+k} \beta_q^{i-k}(\rho)\right).$$

Since, with high probability (which we take to hold with certainty), each of these trapezoids is crossed by at most $\left(\frac{a \log \rho}{\rho}\right)^i m_{\tau_0}$ curves of $C^*_{\tau_0}$, the total number of curve-cell crossings with these trapezoids is at most

$$(ac\beta_q(\rho) \log \rho)^i \left(\rho / \beta_q(\rho)\right)^k m_{\tau_0}.$$

As in the preceding subcase, for both possible values of k, summing over all depths $i = k+1, \ldots, j$, and using the choices of j and ρ, this can be bounded by

$$O\left((ac\beta_q(\rho) \log \rho)^j \rho^k m_{\tau_0}\right) = O\left(h_{\tau_0}^{1/2} (\xi_\tau / s)^\varepsilon m_{\tau_0}\right).$$

Hence the total number of curve-cell crossings within τ_0, taking also into account the case $h_{\tau_0} = 0$, is $O\left(\max\{1, h_{\tau_0}\}^{1/2}(\xi_\tau/s)^\varepsilon m_{\tau_0}\right)$.

We sum the bound just derived over all cells τ_0 of $\mathcal{A}(R'')$, calibrate the value of ε appropriately, and use the facts that the number of cells τ_0 is $O(s)$, that $m_{\tau_0} \le m_\tau/s$, and that $s = \lceil m_\tau^2/X_\tau \rceil$. This yields the following overall bound:

$$\begin{aligned}
O\left(\left(\frac{\xi_\tau}{s}\right)^\varepsilon\right) \cdot \sum_{\tau_0} O(m_{\tau_0} \max\{1, h_{\tau_0}\}^{1/2}) &= \\
O\left(\frac{m_\tau \xi_\tau^\varepsilon}{s^{1+\varepsilon}} \cdot \left(\sum_{\tau_0}(1 + h_{\tau_0})\right)^{1/2} \cdot s^{1/2}\right) &= \\
O\left(m_\tau \xi_\tau^\varepsilon \frac{(h_\tau + s)^{1/2}}{s^{1/2+\varepsilon}}\right) &= \\
O\left(\frac{\xi_\tau^\varepsilon}{s^\varepsilon}\left(m_\tau + \frac{m_\tau h_\tau^{1/2}}{s^{1/2}}\right)\right) &= \\
O\left((X_\tau^{1/2} h_\tau^{1/2} + m_\tau)\xi_\tau^\varepsilon\right)&, \qquad (8)
\end{aligned}$$

for any $\varepsilon > 0$.

If $\xi_\tau < s$ then τ_0 remains intact and the number of crossings between curves and trapezoids within τ_0 is thus m_{τ_0}. We sum this over all $O(\xi_\tau)$ cells τ_0 and use the fact that $m_{\tau_0} \le m_\tau/\xi_\tau$ for each τ_0 to obtain the bound $O(m_\tau)$, which is subsumed in (8).

The case $X_\tau < m_\tau$. So far in the description of the second stage of the decomposition we have assumed that $X_\tau \ge m_\tau$. We now address the case $X_\tau < m_\tau$. By breaking each curve of C_τ^* at the points where it crosses other curves, we obtain a collection of pairwise openly disjoint curves, whose number is only $O(m_\tau)$. Assuming first that $\xi_\tau \le m_\tau$, we now sample each (new) curve in C_τ^* with probability ξ_τ/m_τ, obtaining a random sample R^* of expected size ξ_τ. The expected complexity of the vertical decomposition of $\mathcal{A}(R^*)$ is thus also $O(\xi_\tau)$. By further refining the decomposition, we obtain a collection of $O(\xi_\tau)$ trapezoids, each crossed by at most $O(m_\tau/\xi_\tau)$ curves, for a total of $O(m_\tau)$ crossings between curves and cells. If $m_\tau < \xi_\tau$, we "sample" all curves in C_τ^*, and construct the vertical decomposition of their arrangement. This yields $O(m_\tau) = O(\xi_\tau)$ trapezoids, each of which crosses no curve of C_τ^*.

Completion. We now form the final 2-dimensional decomposition, by taking $\partial\tau^*$ into account. In the description below we address the more involved construction of the case $X_\tau \ge m_\tau$. The derived bounds can be shown to hold also when $X_\tau < m_\tau$ (with a significantly simpler analysis).

The final decomposition in the case $X_\tau \ge m_\tau$ is formed as follows. The hierarchy of trapezoids constructed so far is induced by various samples of (pieces of) curves from C_τ^*. Let Γ_τ denote the collection of all curve portions that constitute the floors and ceilings of all these trapezoids. By construction, no two curve portions in Γ_τ intersect transversally. (Some pairs, constituting, e.g., floors of trapezoids that are nested in the hierarchy, may partially overlap; this has no effect on the analysis about to be presented.) Clearly, the number of trapezoids is $\Theta(|\Gamma_\tau|)$.

Consider now the union Γ'_τ of Γ_τ with the set of arcs forming $\partial\tau^*$. The arcs of Γ'_τ are also pairwise openly disjoint (recalling that the arcs of Γ_τ have been clipped at their points of intersection with $\partial\tau^*$). Form the vertical trapezoidal decomposition of Γ'_τ. Using (7), the number of trapezoids in this decomposition is

$$O(|\Gamma'_\tau|) = O((1 + h_\tau^{1/2})\xi_\tau^{1+\varepsilon} + \xi_\tau + h_\tau) = O((1 + h_\tau^{1/2})\xi_\tau^{1+\varepsilon} + h_\tau).$$

We retain only those trapezoids that are fully contained in τ^* (the others are disjoint from τ^*).

We next consider the number of crossings between the curves of C^*_τ and the new trapezoids. Each such crossing can be charged to a crossing of a curve $\gamma \in C^*_\tau$ with the boundary of a new trapezoid σ (unless γ is fully contained in σ; the number of such latter pairs is clearly at most m_τ). If such a crossing occurs on the floor or ceiling of σ, then either it is also a crossing with the boundary of an old trapezoid, and is thus counted in (8), or it is an endpoint of a curve in C^*_τ (lying on a boundary component of $\partial\tau^*$), and the number of such endpoints is at most $2m_\tau$. If it occurs at a vertical wall erected from an endpoint (or a locally x-extreme point) p of some arc in Γ_τ, then the new wall is equal to or is shorter than the old wall erected from p. Hence the number of such crossings is also upper bounded by (8). The only remaining case is a vertical wall erected from some vertex of $\partial\tau^*$ or from a locally x-extreme point on some arc of $\partial\tau^*$. The number of such walls is $O(\xi_\tau)$, and any such wall is fully contained in an old external trapezoid, and is thus crossed by at most

$$O((a \log \rho/\rho)^j m_{\tau_0}) = O((a \log \rho/\rho)^j (m_\tau/s))$$

curves of C^*_τ. Hence the total number of crossings of this kind is (recall that $\rho^j = \Theta(\xi_\tau/s)$)

$$O(\xi_\tau (a \log \rho/\rho)^j (m_\tau/s)) = O(m_\tau \xi_\tau^\varepsilon),$$

for any $\varepsilon > 0$. This bound also takes care of the case $\xi_\tau < s$, and, as mentioned above, it also trivially holds when $X_\tau < m_\tau$.

The new decomposition is clearly a partition of τ^* into subcells (trapezoids) of constant description complexity. Each of these subcells is lifted vertically in the z-direction to within τ, thereby obtaining a partition of τ itself. The collection of all these partitionings, over all cells τ of $\mathcal{A}_1$, constitutes our final decomposition.

Since each resulting (3-dimensional) cell has constant description complexity, it follows by the ε-net theory of Haussler and Welzl [14] that, with high probability, each of them is crossed by at most $\frac{a'n}{r} \log r$ surfaces of S, for an appropriate absolute constant $a' > 0$, so it is an $O((\log r)/r)$-cutting of S.

Lemma 2.2. *(a) The total number of cells of the above decomposition is $O(r^{3+\varepsilon})$, for any $\varepsilon > 0$.*

(b) The total number of crossings between the curves of C and these cells is $O(mr^{1+\varepsilon})$, for any $\varepsilon > 0$.

Proof: (a) By (7), the number of cells is

$$O\left(\sum_{\tau \in \mathcal{A}_1} \left((1 + h_\tau^{1/2})\xi_\tau^{1+\varepsilon} + h_\tau\right)\right) = O\left(r^2 + \sum_{\tau \in \mathcal{A}_1} (1 + h_\tau^{1/2})\xi_\tau^{1+\varepsilon}\right).$$

We analyze the quantity $O\left(\sum_{\tau\in\mathcal{A}_1}(1+h_\tau^{1/2})\xi_\tau^{1+\varepsilon}\right)$. By Lemma 2.1(a), $\sum_\tau \xi_\tau^{1+\varepsilon} = O(r^{3+\varepsilon})$, for any $\varepsilon > 0$. This bound takes care of all cells for which $h_\tau = 0$. The number of cells with $h_\tau > 0$ is only $O(r^2)$. Moreover, the complexity of a single cell τ of $\mathcal{A}_1$ is only $O(r\beta_q(r))$. Indeed, such a cell has a fixed floor and a fixed ceiling, contained in two respective surfaces σ^-, σ^+ of R. We form a collection of curves, consisting of the xy-projections of (i) the intersections of σ^- and σ^+ with all the remaining surfaces of R, (ii) the silhouettes of the surfaces in R, and (iii) the curves in R'. We obtain a collection of $O(r)$ curves in the plane, and it is easily seen that τ^* is a cell of their arrangement. Hence the complexity of τ^*, and thus of τ, is $O(r\beta_q(r))$, as claimed (see [18] for details). Hence

$$\sum_{\tau\in\mathcal{A}_1} h_\tau^{1/2}\xi_\tau^{1+\varepsilon} = O\left(\left(\sum_{\tau\in\mathcal{A}_1} h_\tau\right)^{1/2} \cdot (r^2)^{1/2}\cdot r^{1+\varepsilon}\right) = O(r^{3+\varepsilon}),$$

for any $\varepsilon > 0$, and this establishes (a).

(b) By (8) and the preceding discussion, the number of crossings is

$$\sum_\tau O\left((X_\tau^{1/2}h_\tau^{1/2} + m_\tau)\xi_\tau^\varepsilon\right),$$

for any $\varepsilon > 0$. Using (1) and (3), the Cauchy-Schwarz inequality, and Lemma 2.1(a,b), and re-calibrating ε, this can be upper bounded by

$$O(r^\varepsilon)\cdot\left[\sum_\tau O(X_\tau^{1/2}h_\tau^{1/2}) + \sum_\tau O(m_\tau)\right] =$$

$$O(r^\varepsilon)\cdot\left(\sum_\tau X_\tau\right)^{1/2}\cdot\left(\sum_\tau h_\tau\right)^{1/2} + O(mr^{1+\varepsilon}) = O(mr^{1+\varepsilon}),$$

for any $\varepsilon > 0$. $\square$

By replacing r by $ar\log r$, for an appropriate absolute constant a, as discussed above, we obtain the following main result:

Theorem 2.3. *Let S be a set of n surfaces in $\mathbb{R}^3$ of constant description complexity, and let C be a set of m curves in $\mathbb{R}^3$ of constant description complexity. Let $1 \le r \le \min\{m,n\}$ be a given parameter. Then there exists a $(1/r)$-cutting Ξ of S of size $O(r^{3+\varepsilon})$, for any $\varepsilon > 0$, such that the number of crossings between the curves of C and the cells of Ξ is $O(mr^{1+\varepsilon})$.*

Remarks. (1) We have ignored so far the algorithmic issue of constructing the cutting. However, the proof is constructive. Moreover, since at each step of the second decomposition stage, we deal with samples of only $O(1)$ curves, the overall cost of the construction can be shown to be $O(nr^{2+\varepsilon} + mr^{1+\varepsilon})$, for any $\varepsilon > 0$. Recall that in the proof we assume that each random sample is a good sample. This can be algorithmically enforced by the standard approach of repeatedly sampling until a good sample is found. Since we only use constant-size samplings in the second decomposition stage, verifying that a sample is good is inexpensive. This approach increases the running time of the algorithm by a constant factor on expectation.

(2) Theorem 2.3 only bounds the overall number of crossings between the curves and cells. A stronger result would be to show that, in addition, each cell of the cutting is crossed by $O(m/r)$ curves of C. We have not carried out this extension, but we believe that this stronger property can be achieved via a modified version of the preceding analysis.

3 The Complexity of a Multiple Zone

Let S and C be as above. Define the zone $Z(C)$ of C in $\mathcal{A}(S)$ to be the collection of all cells of $\mathcal{A}(S)$ that are crossed by at least one curve of C.

Theorem 3.1. *The complexity of $Z(C)$ is $O(m^{1/2}n^{2+\varepsilon})$, for any $\varepsilon > 0$.*

Proof: Since the complexity of the entire arrangement is $O(n^3)$, the bound in the theorem is nontrivial only when $m = O(n^2)$, which is what we assume in the proof. Fix a parameter r, and construct a C-sensitive $(1/r)$-cutting of $\mathcal{A}(S)$, consisting of $O(r^{3+\varepsilon})$ cells, each crossed by at most n/r surfaces of S, so that the total number of crossings between these cells and the curves of C is at most $O(m^{1+\varepsilon}r)$.

Fix a cell τ of the cutting. Let S_τ (resp., C_τ) denote the set of surfaces of S (resp., curves of C) that cross τ, clipped to within τ. The complexity of $Z(C) \cap \tau$ can be upper bounded as follows: First, the zone of a single curve in an arrangement of N surfaces of constant description complexity is $O(N^{2+\varepsilon})$, for any $\varepsilon > 0$ [13]. Hence, the overall complexity of the $|C_\tau|$ separate zones of each of the curves in C_τ in $\mathcal{A}(S_\tau)$ is at most $O(|C_\tau||S_\tau|^{2+\varepsilon})$. In addition, portions of the boundary of the external cell of $\mathcal{A}(S_\tau)$ may also belong to $Z(C)$, because they may bound cells of $\mathcal{A}(S)$ that are crossed by curves of C that do not cross τ. The complexity of this external cell is $O(|S_\tau|^{2+\varepsilon})$. Hence, putting $m_\tau = |C_\tau|$, the overall complexity of $Z(C)$ is (we use the same ε both in the bounds in Theorem 2.3 and for the bound on the complexity of the zone of a curve)

$$O\left(\sum_\tau (m_\tau + 1)\left(\frac{n}{r}\right)^{2+\varepsilon}\right) = O\left(\frac{mn^{2+\varepsilon}}{r} + n^{2+\varepsilon}r\right),$$

where we use Theorem 2.3 to infer that $\sum_\tau m_\tau = O(mr^{1+\varepsilon})$. Choosing $r = m^{1/2}$ completes the proof of the theorem. $\Box$

Remark: A lower bound for $Z(C)$ is $\Omega(m^{2/3}n^{5/3})$. To establish it, take a planar arrangement of $n/2$ lines that has m distinct faces of overall complexity $\Theta(m^{2/3}n^{2/3})$. Lift each of these lines to a vertical plane in three dimensions, and add to the resulting arrangement $n/2$ additional horizontal planes. The resulting collection of n planes is our set S. For the set C of curves, take m vertical lines, each intersecting the xy-plane at a point inside one of the m marked faces. The complexity of the multiple zone $Z(C)$ is easily seen to be $\Theta(m^{2/3}n^{5/3})$.

Acknowledgments

The authors would like to thank Gábor Tardos for suggesting the problem, and thank him and János Pach for valuable discussions. We also deeply thank two anonymous referees for very helpful comments on the paper, which have led to considerable improvements. In addition, we are greatly indebted to one of the referees, who has pointed out a serious technical problem in an earlier version of the paper.

References

[1] P. K. Agarwal, J. Matoušek and O. Schwarzkopf, Computing many faces in arrangements of lines and segments, *SIAM J. Comput.* 27:491–505, 1998.

[2] P. K. Agarwal and M. Sharir, Arrangements and their applications, In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 49–119, Elsevier Science Publishers B.V., North-Holland, Amsterdam, 2000.

[3] B. Aronov, M. Pellegrini, and M. Sharir, On the zone of a surface in a hyperplane arrangement, *Discrete Comput. Geom.* 9:177–186, 1993.

[4] M. de Berg, L. J. Guibas, and D. Halperin, Vertical decompositions for triangles in 3-space, *Discrete Comput. Geom.*, 15:35–61, 1996.

[5] M. de Berg and O. Schwarzkopf, Cuttings and applications, *Internat. J. Comput. Geom. Appls.*, 5:343–355, 1995.

[6] B. Chazelle, Cutting hyperplanes for divide-and-conquer, *Discrete Comput. Geom.* 9:145–158, 1993.

[7] B. Chazelle and J. Friedman, A deterministic view of random sampling and its use in geometry, *Combinatorica*, 10:229–249, 1990.

[8] K. Clarkson, H. Edelsbrunner, L. J. Guibas, M. Sharir and E. Welzl, Combinatorial complexity bounds for arrangements of curves and spheres, *Discrete Comput. Geom.* 5:99–160, 1990.

[9] K. Clarkson, New applications of random sampling in computational geometry, *Discrete Comput. Geom.* 2:195–222, 1987.

[10] K. Clarkson and P. Shor, Applications of random sampling in computational geometry, II, *Discrete and Computational Geometry*, 4:387–421, 1989.

[11] D.P. Dobkin and D.G. Kirkpatrick, Fast detection of polyhedral intersection. *Proc. 9th Int. Colloq. Automata, Languages, and Programming*, Springer-Verlag LNCS 140, 154–165, 1982.

[12] E. Ezra and M. Sharir, Counting and representing intersections among triangles in three dimensions, *Proc. 20th ACM Symposium on Computational Geometry*, 210–219, 2004.

[13] D. Halperin and M. Sharir, Almost tight upper bounds for the single cell and zone problems in three dimensions, *Discrete Comput. Geom.* 14:385–410, 1995.

[14] D. Haussler and E. Welzl, Epsilon-nets and simplex range queries, *Discrete Comput. Geom.*, 2:127–151, 1987.

[15] V. Koltun and M. Sharir, Curve-sensitive cuttings, *Proc. 19th ACM Symposium on Computational Geometry*, 136–143, 2003.

[16] J. Matoušek, Construction of ϵ-nets, *Discrete Comput. Geom.*, 5:427–448, 1990.

[17] J. Matoušek. Range searching with efficient hierarchical cuttings, *Discrete Comput. Geom.*, 10:157–182, 1993.

[18] M. Sharir and P.K. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.